

# Styleformer: Transformer based Generative Adversarial Networks with Style Vector

Jeeseung Park\*  
mAy-I Inc.

jspark@may-i.io

Younggeun Kim\*  
MINDsLab Inc.

younggeun@minds-lab.ai

## Abstract

We propose *Styleformer*, a generator that synthesizes image using style vectors based on the Transformer structure. In this paper, we effectively apply the modified Transformer structure (e.g., Increased multi-head attention and Pre-layer normalization) and introduce novel Attention Style Injection module which is style modulation and demodulation method for self-attention operation. The new generator components have strengths in CNN's shortcomings, handling long-range dependency and understanding global structure of objects. We present two methods to generate high-resolution images using *Styleformer*. First, we apply Linformer in the field of visual synthesis (*Styleformer-L*), enabling *Styleformer* to generate higher resolution images and result in improvements in terms of computation cost and performance. This is the first case using Linformer to image generation. Second, we combine *Styleformer* and StyleGAN2 (*Styleformer-C*) to generate high-resolution compositional scene efficiently, which *Styleformer* captures long-range dependencies between components. With these adaptations, *Styleformer* achieves comparable performances to state-of-the-art in both single and multi-object datasets. Furthermore, groundbreaking results from style mixing and attention map visualization demonstrate the advantages and efficiency of our model.

## 1. Introduction

Generative Adversarial Network (GAN) [21] is one of the widely used generative model. Since the appear of DCGAN [43], convolution operations have been considered essential for high-resolution image generation and stable training. Convolution operations are created under the assumption of the locality and stationarity of the image (i.e., inductive bias), which is advantageous for image processing [37]. Through convolution neural networks (CNNs) with this strong inductive bias, GAN have efficiently gen-

erated realistic, high-fidelity images.

However, drawbacks of CNNs clearly exist. Local receptive field of CNNs makes model difficult to capture long-range dependency and understanding global structure of object. Stacking multiple layers can solve this problem, but this leads to another problem of losing spatial information and fine details [55]. Moreover, sharing kernel weights across locations leads to unstable training when the pattern or styles differ by location in the image [56]. This is also related to the poor quality of generated structured images or compositional scenes (e.g., outdoor scenes), unlike the generation of a single object (e.g., faces)

In this paper, we propose *Styleformer*, a generator that uses style vectors based on the Transformer structure. Unlike CNNs, *Styleformer* utilizes self-attention operation to capture long-range dependency and understand global structure of objects efficiently. Furthermore, we overcome computation problem of Transformer and show superior performance not only in low-resolution but also in high resolution images. Specifically, we introduce the following three models:

1) *Styleformer* - The basic block of *Styleformer* is based on Transformer encoder, so we introduce components that need to be changed for stable learning. Inspired by MobileStyleGAN [3], we enhance the multi-head attention in original Transformer by increasing the number of heads, allowing model to generate image efficiently. We also modify layer normalization, residual connection, and feed-forward network (Section 3.2). Moreover, we introduce novel *attention style injection* module, suitable style modulation, and demodulation method for self-attention operation (Section 3.3). This design allows *Styleformer* to generate image stably, and enables model to handle long-range dependency and understand global structures.

2) *Styleformer-L* - We sidestep scalability limitation arising from the quadratic mode of attention operation by applying Linformer [50] (*Styleformer-L*). As such, *Styleformer-L* can generate high-resolution images with linear computational costs. This paper is the first case to apply Linformer in the field of visual synthesis (Section 3.4).

\*Equal contribution. The order was determined randomly.

3) *Styleformer-C* - We further combine Styleformer and StyleGAN2, applying Styleformer at low resolution and style block of StyleGAN2 at high resolution (Styleformer-C). As can be seen from our experiments and analysis (e.g., style mixing and visualizing attention map), we show that Styleformer-C with the structure above can generate compositional scenes efficiently, and showing flexibility of our model. In detail, we prove that Styleformer in low resolution help model to capture long-range dependency between components, and style block in high resolution help model to refine the details of each components such as color or texture. This novel blending structure enables fast training, which is the advantage of StyleGAN2, while maintaining the advantages of Styleformer that can generate structured images.(Section 4).

Styleformer achieves comparable performances to state-of-the-art in both single and multi-object datasets. We record FID 2.82 and IS 10.00 at the unconditional setting on CIFAR-10. These results outperform all GAN-based models including StyleGAN2-ADA [32] which recently recorded state-of-the-art. As can be expected, Styleformer show strength especially in multi-object images or compositional scenes generation (e.g., CLEVR, Cityscapes). Styleformer-C records FID 11.67, IS 2.27 in CLEVR, and FID 5.99, IS 2.56 in Cityscapes, showing better performance than pure StyleGAN2.

## 2. Related Work

After origion of GAN [21], various methods [2, 31, 41, 42] have been proposed to enhance its training stability and performance. As a result, fidelity and diversity of the generated images have dramatically improved. In addition to image synthesis task, GAN has been widely adopted in various tasks, such as image-to-image translation [27, 58], super resolution [38], image editing [55], and style transfer [7]. In particular, StyleGAN-based architectures have been applied for various applications [16, 59, 60]. However, since all of these models are based on convolution backbones, they have met with only limited success on generating complex or compositional scenes [29].

Transformer [49] was first introduced to the natural language processing(NLP) domain, achieving a significant advance in NLP. Recently, there were efforts to utilize Transformer in the computer vision field [4, 12, 57]. Using huge amounts of data and a transformer module, ViT [12] obtains comparable result with state-of-the-art model in the existing CNN based image classification model [35, 47]. Inspired by [12], various models such as [22, 39, 53] emerges based on this structure. There have also been attempts to utilize transformer for tasks such as video understanding [4] and segmentation [57] as well as image classification. Even in GAN, there have been attempts to utilize Transformer: GANformer [26] proposes a bipartite Transformer struc-



Figure 1. High-resolution compositional scenes generated by Styleformer.

ture and applies it to StyleGAN [33, 34]. With this structure, GANformer successfully advance the generative modeling of structured images and scenes, which have been challenging in existing GAN. However, they use a bipartite attention, differ from the self-attention operation. TransGAN [28] demonstrates a convolution-free generator based on the structure of vanilla GAN, which doesn't show good performance compared to state-of-the-art model.

Unlike these studies, Styleformer generate images with self-attention operation using style vector and showing comparable performance state-of-the-art models [33, 34]. Previous methods (TransGAN) mainly use pre-defined sparse attention patterns for efficient attention mechanism, but we explore the low-rank property in self-attention. Our model can generate high resolution images ( $512 \times 512$ ) with reduced computation complexity, while GANformer and TransGAN show a maximum of  $256 \times 256$  image synthesis.

## 3. Styleformer

### 3.1. Styleformer Architecture

Figure 2a shows the overall architecture of Styleformer, and in Figure 2b we show Styleformer encoder network, the basic block of Styleformer. Like existing synthesis network of StyleGAN, our generator is conditioned on a learnable constant input. The difference is that the constant input ( $8 \times 8$ ) is flattened (64) to enter the Transformer-based encoder. Then the input which is combined with learnable positional encoding passes through the Styleformer encoder. Styleformer encoder is based on Transformer encoder, but there are several changes to generate an image efficiently, which will be discussed in Section 3.2.

After passing several encoder blocks in each resolution, we proceed bilinear upsample operation by reshaping encoder output to the form of a square feature map. After upsampling, flatten process is carried out again to match the input form of the Styleformer encoder. This process is repeated until the feature map resolution reaches the target

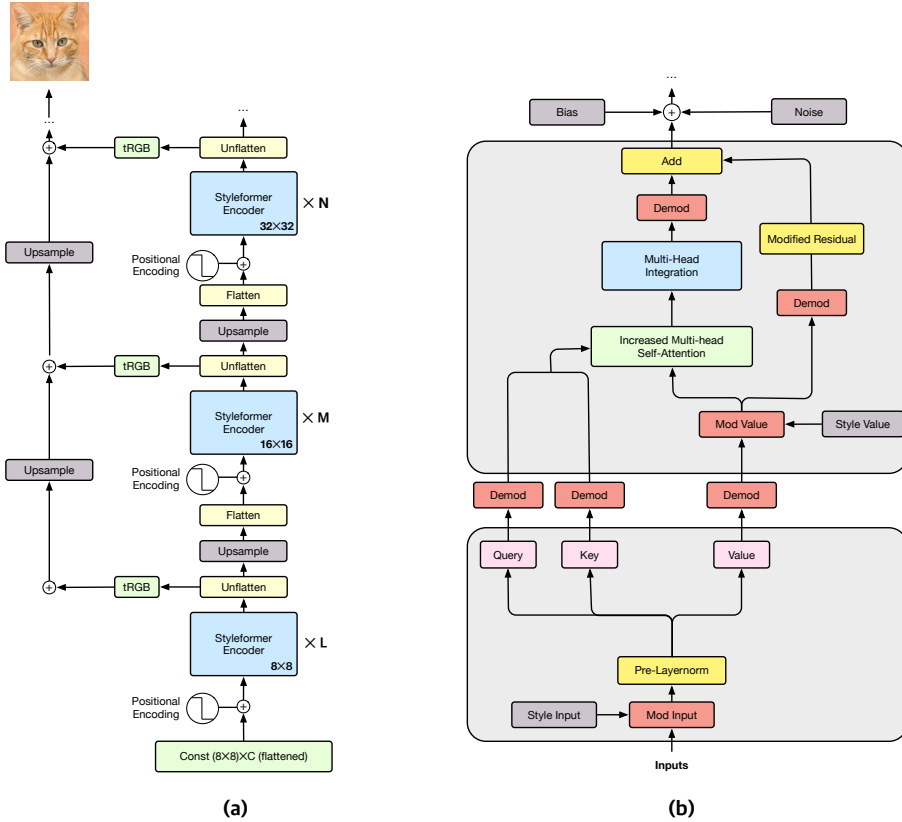


Figure 2. (a) Overall Architecture of Styleformer. (b) Styleformer encoder structure, which is the basic block of Styleformer.

image resolution. For each resolution, the number of the Styleformer encoder and hidden dimension size can be chosen as hyperparameters.

### 3.2. Styleformer Components from Transformer

**Increased Multi-Head Attention** Modern vision architectures allow communications between different channels and different spatial locations (i.e., pixels) [48]. Conventional CNNs perform the above two communications at once, but these communications can be clearly separated like depthwise separable convolutions [24]. We also separate the pixel-communication (self-attention), channel-communication operations (multi-head integration) in the Transformer encoder. However, in depthwise separable convolutions, distinct convolution kernels are applied to each channel, unlike the self-attention operation share only one huge kernel  $A$  (i.e., attention map). With same kernel applied to each channel, diversity in generated image can be decreased.

We overcome this problem by increasing the number of heads of multi-head attention (Increased multi-head attention). Then, the created attention map will be different for each head, and so the kernel applying operation. Then the attention maps will be created for each head, making the

channels in each head meet different kernels. However, increasing the number of heads too much may cause attention map to not be properly created, resulting in poor performance. We demonstrate experimentally that increasing the number of heads improves performance only when the depth is at least 32, as shown in Figure 3. Therefore, we fix the depth to 32 for all future experiment. More details about increased multi-head attention can be found in Appendix C.

**Pre-Layer Normalization** We change the position of layer normalization in Transformer encoder. The layer normalization of the existing Transformer comes after a linear layer that integrates multi-heads (Post-Layer normalization). We hypothesis that the role of layer normalization in a Transformer is the preparation of generating an attention map. If we perform layer normalization at the end of Styleformer encoder (*Layernorm B* in Figure 4), style modulation is applied before making query and key, which can disturb learning attention map. This is supported by ablation study and attention map analysis in Table 1 and Appendix B, respectively. Therefore, to solve this problem, we proceed layer normalization before operation making query, key and value (*Pre-Layernorm* in Figure 2b)

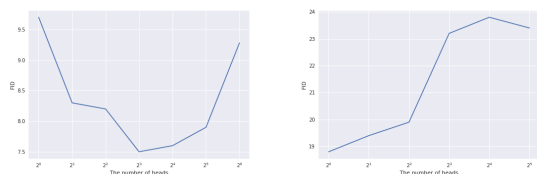


Figure 3. It shows FID on CIFAR-10 with one layer Styleformer, which hidden dimension size is fixed as 256 and 32, respectively. Both experiments show the best result when the depth is 32.

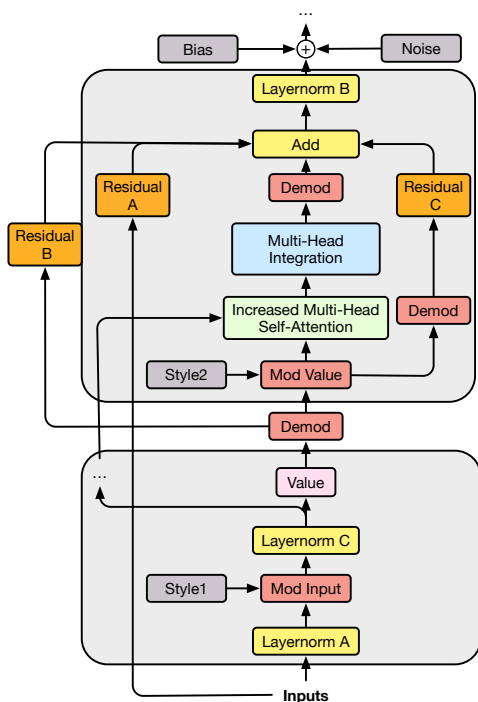


Figure 4. Styleformer encoder structure for ablation study, including residual connection, layer normalization, attention style injection.

**Modified Residual Connection** Unlike the Transformer encoder, input feature map is scaled by style vector (*Mod input* in Figure 2b) in Styleformer encoder. We hence find the residual connection suitable for scaled input. After ablation study, we apply residual connection like *Modified Residual* in Figure 2b. Demodulation operation is additionally performed in residual connection, which will be described in Section 3.3. Table 1 presents ablation details of residual connection.

**Eliminating Feed-Forward Network** As can be seen Table 1, we remove the feed-forward structure because eliminating feed forward structure makes the model perform better and more efficient.

### 3.3. Attention Style Injection

Unlike vanilla GAN, StyleGAN generates an image with layer-wise style vectors as inputs, enabling controllable generation via style vectors, i.e., scale-specific control. Specifically, style vector scales the input feature map for each layer, i.e., style modulation, amplifying certain feature maps. For scale-specific control, this amplified effect must be removed before entering the next layer. StyleGAN allows scale-specific control through a normalization operation called AdaIN operation [13, 14, 18, 25], which normalizes each feature map separately, then scale and bias each feature map with style vector. StyleGAN2 is an advanced form of StyleGAN and addresses the artifact problem caused by the AdaIN operation, solving it by demodulation operation. While the AdaIN operation normalizes the output feature map directly, demodulation operation is based on statistical assumptions about the input feature map. For details, similar to the goal of normalization operation, demodulation operation aims to have an output feature map with a unit standard deviation while assuming that the input feature maps have a unit standard deviation, i.e., statistical assumption. Our goal is to design a Transformer-based generator that generates images through style vector while enabling scale specific control. Therefore, we propose style modulation, demodulation method for the self-attention operation, i.e., *Attention style injection*.

**Modulation for Self-Attention** Just as the input feature map is scaled by style vector in the style block of StyleGAN2, the input feature map in the Styleformer encoder is also scaled by style vector (*Mod Input* in Figure 2b). But unlike convolution operation in StyleGAN2, there are two steps in self-attention operation: dot product of query and key to create an attention map (i.e. kernel), weighted sum of the value with calculated attention map. We hypothesis that the style vector applied to the operation in each step should be different. Therefore, we perform style modulation twice as in Figure 2b (*Mod Input*, *Mod Value*). This hypothesis is supported in Table 1. In Figure 2b, *Style Input* is a style vector for input, and *Style Value* is a style vector only for value. Two style vectors are created through common mapping networks as in StyleGAN but different learned affine transformations.

**Demodulation for Query, Key, Value** As shown in Figure 2b, Styleformer encoder creates query ( $Q$ ), key ( $K$ ), and value ( $V$ ) through linear operation to the input feature map scaled with *Style Input* vector. After that,  $V$  will be modulated with *Style Value* vector additionally, so the demodulation operation for removing scaled effect of *Style Input* is clearly required. Also, we observe that when an attention map is created with  $Q$ ,  $K$  from input scaled by *Style In-*

Method	Style1	Style2	Style1=Style2	Residual A	Residual B	Residual C	Layernorm A	Layernorm B	Layernorm C	Feed-Forward	FID
<b>Baseline</b>	O	O	X	X	X	O	X	X	O	X	8.56
<b>Attention Style Injection</b>	O	X	-	X	X	O	X	X	O	X	11.01
	X	O	-	X	X	O	X	X	O	X	11.40
	O	O	O	X	X	O	X	X	O	X	10.27
<b>Residual Connection</b>	O	O	X	X	X	X	X	X	O	X	19.09
	O	O	X	O	X	X	X	X	O	X	14.70
	O	O	X	X	O	X	X	X	O	X	9.94
<b>Layer Normalization</b>	O	O	X	X	X	O	O	X	X	X	9.00
	O	O	X	X	X	O	X	O	X	X	10.96
<b>Feed-Forward</b>	O	O	X	X	X	O	X	X	O	O	14.75

Table 1. Ablation details of Styleformer components. Ablation study was conducted using small version of Styleformer with CIFAR-10 dataset, trained for 20M images. See Appendix A for further implementation details.

put, specific value in the attention map becomes very large, demonstrated in Appendix B. This prevents the attention operation from working properly. We sidestep this problem with demodulation operation to  $Q$ ,  $K$ , before creating attention map. Eventually, demodulation operation is all required for  $Q$ ,  $K$ , and  $V$ .

Let’s first look at the style modulation to the input, i.e., *Mod Input*. Each flattened input feature map is scaled through a style vector, which is equivalent to scaling the linear weight:

$$w'_{ij} = s_i \cdot w_{ij}, \quad (1)$$

where  $w$  is original linear weight to make  $(Q, K, V)$  from flattened input feature map, and  $w'$  is modulated linear weight.  $s_i$  is  $i$ th component of style vector, which scales  $i$ th flattened input feature map, and  $j$  means the dimension of  $(Q, K, V)$ . Assuming that flattened input feature maps have unit standard deviation (i.e., statistical assumption of demodulation), after passing style modulation and linear operation, a standard deviation of output is as follows:

$$\sigma_j = \sqrt{\sum_i w'_{ij}{}^2}. \quad (2)$$

We scale output activations for each dimension of  $Q$ ,  $K$ , and  $V$  by  $1/\sigma_j$  (i.e., demodulation), making  $Q$ ,  $K$ , and  $V$  back to unit standard deviation.

**Demodulation for Encoder Output** After demodulation operation to  $Q$ ,  $K$ , and  $V$ , Styleformer encoder performs style modulation to  $V$  (*Mod Value*), weighted sum of  $V$  with attention map (*Increased Multi-head Self-attention*), and then performs linear operation (*Multi-Head Integration*), as shown in Figure 2b. Encoder output will be input for next encoder, so demodulation operation is necessary. We show in Appendix D that, assuming  $V$  has a unit standard deviation (This can be assumed because of demodulation for  $V$ ), the standard deviation of Styleformer encoder

output can be derived as follows:

$$\sigma'_{lk} = \sqrt{\sum_l A_l \cdot \sum_j w'_{jk}{}^2}, \quad (3)$$

where  $w'_{jk} = s_j \cdot w_{jk}$ , i.e., modulated linear weight.  $s_j$  scales  $j$ th feature map of  $V$ , and  $k$  enumerates the flattened output feature map. Attention map  $A$  is computed same as existing Transformer: dot products of  $Q$  and  $K$ , divide each by square root of depth, and softmax function.  $A_l$  denotes attention score vector for  $l$ th pixel.

However, there are two problems with demodulation by simply scaling each flattened output feature map  $k$  with  $1/\sigma'_{lk}$  (Equation 3). First, scaling output feature map  $k$  with  $1/\sigma'_{lk}$  will normalize each pixel as a unit, different from AdaIN operation which normalizes each feature map as a unit. Second, the attention map, which is a matrix derived from  $Q$  and  $K$ , is dependent on the input. With input dependent variables, demodulation operations based on statistical assumptions can not be applied as in [19]. Therefore we scale the flattened output feature map  $k$  with  $1/\sigma''_k$  where  $\sigma''_k = \sqrt{\sum_j w'_{jk}{}^2}$ , normalizing each feature map as a unit, and excluding input dependent variables  $A_l$ . Then the standard deviation of output activations will be

$$\sigma_{lk} = \frac{\sigma'_{lk}}{\sigma''_k} = \sqrt{\sum_l A_l}. \quad (4)$$

However in this way, standard deviation of output is not unit, rather approaching to zero when the numbers of pixels increase, as detailed in Appendix D. To prevent this effect, we have applied modified residual connection like *Modified Residual* in Figure 2b. More specifically, we perform linear operation to *Mod Value*, then perform demodulation operation (same as demodulation for query, key, value). With these modulation and demodulation operations in residual connection, variables with unit standard deviation are added to the output. Therefore it helps to keep the final output ac-

tivation having unit standard deviation, when  $\sigma_{jk}$  is close to zero.

### 3.4. High Resolution Synthesis with Styleformer

The main problem in applying Transformer to image generation is the efficiency problem with image resolution. In this section, we introduce two different techniques in Styleformer that can generate high resolution images. We show a method of applying Linformer, making computation complexity to linear. Then, we introduce a method of combining Styleformer and StyleGAN2, which can obtain the advantages of both models.

**Applying Linformer to Styleformer** For high-resolution images, input sequence length of the Styleformer encoder increases quadratically, and the standard self-attention mechanism requires a complexity of  $O(n^2)$  with respect to the sequence length. It means attending to all pixels for each layer is almost impossible for high-resolution image generation. Therefore, we apply Linformer [50] to our model, which projects key and value to the  $k$  dimension when applying self-attention, reducing the time and space complexity from  $O(n^2)$  to  $O(nk)$ . We fix  $k$  to 256 and apply Linformer to the encoder block above  $32 \times 32$  resolution, only when  $n$  is 1024 or higher. We call this model as *Styleformer-L*.

[50] explains that this new self-attention mechanism succeeds because the attention map matrix is low-rank. We observe this can be applied equally to the attention map matrix in the image: in the case of images, the pixel that needs to attend is often in a particular location, not all pixels (e.g. where objects are located in the image), which results in low-rank attention map matrix. Applying Linformer creates a more dense attention map, and also reduces computation. This is proved by spectrum analysis of attention map in Section 4.2. See Appendix E for more details about Styleformer-L.

**Combining Styleformer and StyleGAN2** Even with applying Linformer, it is difficult to generate an image for extremely high resolution like  $512 \times 512$  using only Transformer. We solve this problem by combining Styleformer and StyleGAN2 to generate a high-resolution image, and we call this model *Styleformer-C*. Styleformer-C is composed of Styleformer at low resolution, and style block of StyleGAN2 at high resolution. As demonstrated in 4.1, Styleformer encoder in low resolution help model to capture long-range dependency between components or global shape of object, and style block in high resolution help model to refine the details of each components or objects. In other words, model can capture global interactions efficiently using Styleformer only at low resolution, which

leads to fast training speed. The overall architecture and details of Styleformer-C are described in Appendix F.

## 4. Experiments

We only change the architecture of the generator in StyleGAN2-ADA, i.e., synthesis network, while maintaining the discriminator architecture and loss function. We use Fréchet Inception Distance (FID) [23] and Inception Score (IS) [44], evaluation metrics mainly used in the field of image generation. We compare our model with top GAN models such as StyleGAN2-ADA [32], and model related to our research such as TransGAN. In Section 4.1, we show performance results of Styleformer in low-resolution dataset. Section 4.2 provide evidence for a successful application of Linformer, including performance of Styleformer-L. In Section 4.3, we show high performance of Styleformer-C and prove the advantage and efficiency of our model by style mixing, and attention map visualization.

### 4.1. Low-Resolution Synthesis with Styleformer

Styleformer achieves comparable performance to state-of-the-art in various low-resolution single-object datasets, including CIFAR-10 ( $32 \times 32$ ) [36], STL-10 ( $48 \times 48$ ) [9], and CelebA ( $64 \times 64$ ) [40].

As shown in Table 2, Styleformer outperforms prior GAN-based models, in terms of FID and IS. Especially in CIFAR-10, Styleformer records FID 2.82, and IS 10.00, which is comparable to current state-of-the-art and outperforming StyleGAN2-ADA-tuning. These results indicates that the Styleformer encoder has been modified to generate image successfully. Implementation details are in Appendix A.

### 4.2. Applying Linformer to Styleformer

We experiment our method at Section 3.4 which applies Linformer to Styleformer (Styleformer-L) on CelebA,  $64 \times 64$  resolution, and LSUN-Church [54] dataset resized to  $128 \times 128$  resolution. As shown in Table 3, we find significant improvements in speed and memory and better performance than conventional Styleformer on CelebA. Memory performance is approximately three times more effective and speed performance is 1.3 times better in Styleformer-L. We also succeed in generating images of  $128 \times 128$  resolution with the LSUN-Church dataset, which is difficult with pure Styleformer due to expensive memory.

In addition, in the CelebA dataset, Styleformer-L shows higher performance in terms of FID than Styleformer, improving FID scores from 3.92 to 3.36. To analyze this phenomenon, we extract an attention map from Styleformer for generated CelebA images. As in [50], we apply singular value decomposition into attention map matrix, and plot the normalized cumulative singular value averaged over 1k generated images. As shown in Figure 6, most of the informa-

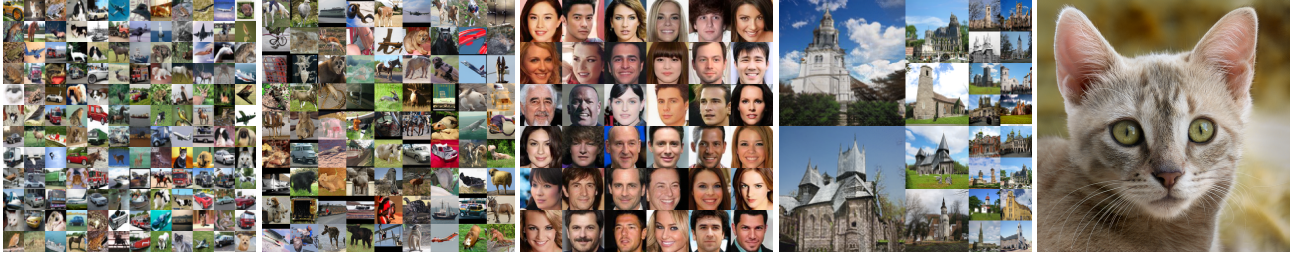


Figure 5. From the left, results generated by Styleformer on CIFAR-10 and STL-10, results generated by Styleformer-L on CelebA and LSUN-church and results generated by Styleformer-C on AFHQ-Cat. For more generated samples, please see Appendix G.

Cifar-10			STL-10			CelebA	
Method	FID ↓	IS ↑	Method	FID ↓	IS ↑	Method	FID ↓
Progressive-GAN [31]	15.52	8.80 ± 0.05	SN-GAN [42]	40.1	9.16 ± 0.12	PAE [5]	49.2
AutoGAN [20]	12.42	8.55 ± 0.10	Improving MMD-GAN [51]	37.64	9.23 ± 0.08	BEGAN-CS [6]	34.14
StyleGAN V2 [34]	11.07	9.18	AutoGAN [20]	31.01	9.16 ± 0.12	PeerGAN [52]	13.95
Adversarial NAS-GAN [20]	10.87	8.74 ± 0.07	Adversarial NAS-GAN [17]	26.98	9.63 ± 0.19	TransGAN-XL [28]	12.23
TransGAN-XL [28]	9.26	9.02 ± 0.11	TransGAN-XL [28]	18.28	10.43 ± 0.17	HDCGAN [11]	8.77
StyleGAN2-ADA [32]	2.92	9.83 ± 0.04	SNGAN-DCD [46]	17.68	9.33	NCP-VAE [1]	5.25
<b>Styleformer</b>	<b>2.82</b>	<b>10.00 ± 0.12</b>	<b>Styleformer</b>	<b>15.17</b>	<b>11.01 ± 0.15</b>	<b>Styleformer</b>	<b>3.92</b>

Table 2. Comparison results between Styleformer and other GAN models on low-resolution datasets. Results of other GAN models are collected from papers that reports their best results. We compute FID, IS in the same way as StyleGAN2-ADA, generating 50k images and compare their statistics against the 50k images from the training set for FID, computing the mean over 10 dependent trials using 5k generated images per trial for IS.

Dataset	Model	FID ↓	Memory per GPU ↓	Speed ↓
CelebA	Styleformer	3.92	14668MiB	6.46
	<b>Styleformer-L</b>	<b>3.36</b>	5316MiB	4.93
LSUN church	Styleformer	-	OOM	-
	<b>Styleformer-L</b>	<b>7.99</b>	8118MiB	9.81

Table 3. Results on Styleformer-L which applies Linformer. “Memory” is measured on 4 Titan-RTX with 16 batch size per GPU and “Speed” means seconds for processing 1k images (sec/1kimg). We use the same hidden dimension and the number of layers in Styleformer and Styleformer-L.

Method	CLEVR		Cityscapes	
	FID ↓	IS ↑	FID ↓	IS ↑
GAN [21]	25.02	2.17	11.57	1.63
k-GAN [45]	28.09	2.21	51.08	1.66
SAGAN [56]	26.04	2.17	12.81	1.68
StyleGAN2 [34]	16.05	2.15	8.35	1.70
VQGAN [15]	32.60	2.03	173.80	<b>2.82</b>
Styleformer-C	<b>11.67</b>	<b>2.27</b>	<b>5.99</b>	2.56

Table 4. Comparison between popular CNN based GAN models and Styleformer-C on CLEVR and Cityscapes. We use the results in [26] for the performance of other models.

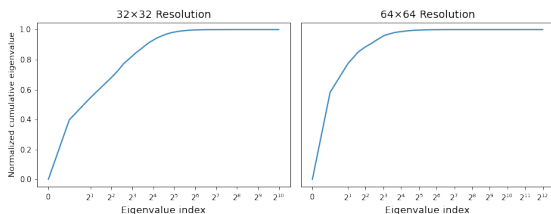


Figure 6. Spectrum analysis of attention map matrix at 32, 64 resolution. We use pretrained Styleformer with CelebA dataset.

tion in the attention map matrix can be recovered from the few large singular value, which means that the rank of attention map matrix is low. With low rank attention map,

Linformer can be applied more efficiently [50].

Therefore, we show the possibility that when applying a self-attention operation for high-resolution images, it is not necessary to apply attention to all pixels and provide scalability to generate high-resolution images using Styleformer-L. See Appendix E for implementation details.

### 4.3. Styleformer can Capture Global Interaction

We experiment our method at Section 3.4 which combines Styleformer and StyleGAN2 (Styleformer-C) on CLEVR(256 × 256) [30] and Cityscapes (256 × 256) [10] for multi-object images and compositional scenes, AFHQ CAT (512 × 512) [8] for high-resolution single-object images.

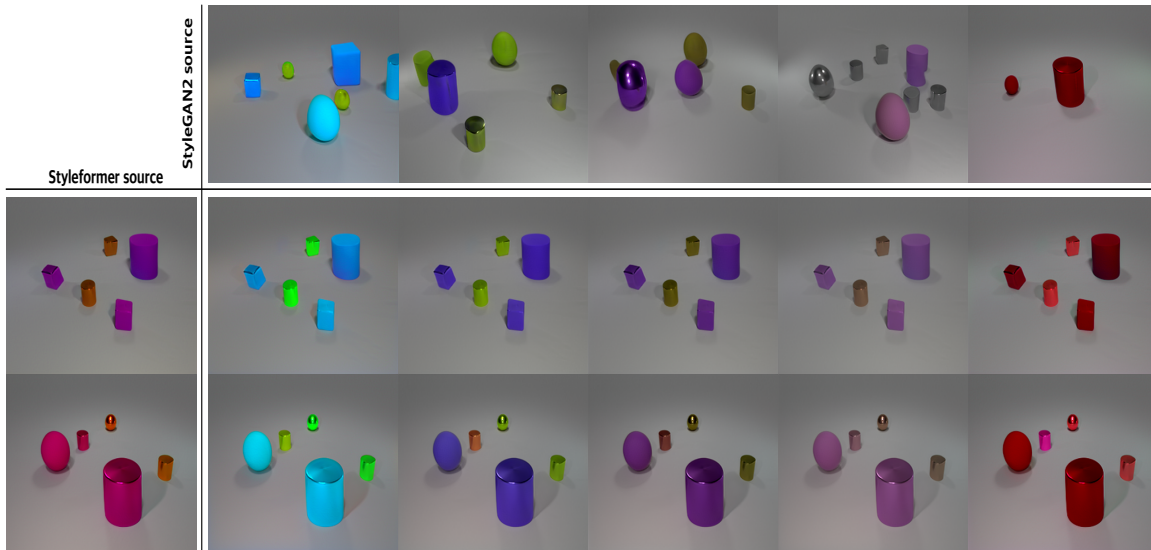


Figure 7. Style mixing experiment with Styleformer-C on CLEVR dataset. The images on the x-axis and y-axis were generated from their respective latent codes (StyleGAN2 source and Styleformer source, respectively); the rest of the images were generated by applying styles from Styleformer source to Styleformer at low resolution and applying styles from StyleGAN2 source to StyleGAN2 at high resolution [33].

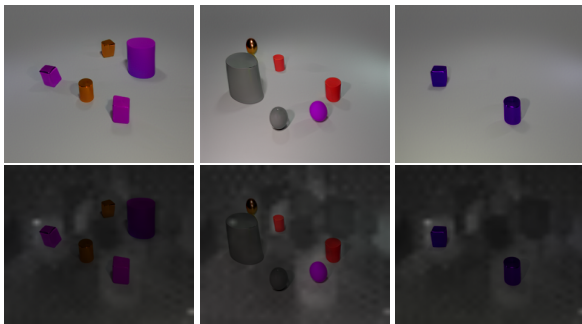


Figure 8. Visualizing attention map in generated CLEVR images.

As shown in Table 4, Styleformer-C records FID 11.67, IS 2.27 in CLEVR, and FID 5.99, IS 2.56 in Cityscapes which is comparable performance to current state-of-the-art, and showing better performance than StyleGAN2 in multi-object images and compositional scenes. This indirectly shows that Styleformer helps model to handle long-range dependency between components.

To show more solid evidence that Styleformer captures global interaction, We conduct style mixing [33] in Styleformer-C. In detail, when generating new image from CLEVR dataset, we use two different latent codes  $z_1, z_2$  and applying  $z_1$  to Styleformer at low resolution and  $z_2$  to StyleGAN2 at high resolution. As shown in Figure 7, style corresponding to Styleformer (low-resolution) brings the basis for structural generation such as the location and structure of objects, while all colors or textures remain same. On the contrary, style corresponding to StyleGAN2 (high-

resolution) brings the color and texture change, while maintaining location and shape of objects. This results directly prove that Styleformer controls global structure between objects, and handles long-range dependency.

In addition, we visualize the attention map to provide more insight into the model’s generating process. Figure 8 shows the concentration of attention to the position where the object exists. These visualizations show that the self-attention operation worked efficiently, enabling the model to perform long-range interaction, overcome the shortcoming of convolution operation.

## 5. Conclusion

We propose Styleformer, a Transformer-based generative network that is novel and effective. We propose a method to efficiently generate images with self attention operation and achieve SOTA performance on various datasets. Furthermore, we propose Styleformer-L, which reduces the complex computation to linear, enabling to generate high-resolution images. We also present a method of efficiently generating a compositional scene while capturing with long-range dependency through Styleformer-C. There still seems to be room for improvement, such as reducing computation cost, but we hope that our work will speed up the application of Transformers to the field of computer vision, helping the development of the computer vision field. However, development of the generative model can create fake media data using synthesized face images (e.g. deep-fake), so particular attention should be paid in the future.



## References

- [1] Jyoti Aneja, Alexander Schwing, Jan Kautz, and Arash Vahdat. Ncp-vae: Variational autoencoders with noise contrastive priors, 2020. 7
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan, 2017. 2
- [3] Sergei Belousov. Mobilestylegan: A lightweight convolutional neural network for high-fidelity image synthesis, 2021. 1
- [4] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding?, 2021. 2
- [5] Vanessa Böhm and Uroš Seljak. Probabilistic auto-encoder, 2020. 7
- [6] Chia-Che Chang, Chieh Hubert Lin, Che-Rung Lee, Da-Cheng Juan, Wei Wei, and Hwann-Tzong Chen. Escaping from collapsing modes in a constrained space, 2018. 7
- [7] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation, 2018. 2
- [8] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains, 2020. 7
- [9] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In Geoffrey Gordon, David Dunson, and Miroslav Dudík, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 215–223, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR. 6
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding, 2016. 7
- [11] J. D. Curtó, I. C. Zarza, Fernando de la Torre, Irwin King, and Michael R. Lyu. High-resolution deep convolutional generative adversarial networks, 2020. 7
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2020. 2
- [13] Vincent Dumoulin, Ethan Perez, Nathan Schucher, Florian Strub, Harm de Vries, Aaron Courville, and Yoshua Bengio. Feature-wise transformations. *Distill*, 2018. <https://distill.pub/2018/feature-wise-transformations>. 4
- [14] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style, 2017. 4
- [15] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis, 2021. 7
- [16] Aviv Gabbay and Yedid Hoshen. Style generator inversion for image enhancement and animation, 2019. 2
- [17] Chen Gao, Yunpeng Chen, Si Liu, Zhenxiong Tan, and Shuicheng Yan. Adversarialnas: Adversarial neural architecture search for gans, 2020. 7
- [18] Golnaz Ghiasi, Honglak Lee, Manjunath Kudlur, Vincent Dumoulin, and Jonathon Shlens. Exploring the structure of a real-time, arbitrary neural artistic stylization network, 2017. 4
- [19] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In Yee Whye Teh and Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR. 5
- [20] Xinyu Gong, Shiyu Chang, Yifan Jiang, and Zhangyang Wang. Autogan: Neural architecture search for generative adversarial networks, 2019. 7
- [21] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. 1, 2, 7
- [22] Ben Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference, 2021. 2
- [23] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium, 2018. 6
- [24] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017. 3
- [25] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization, 2017. 4
- [26] Drew A. Hudson and C. Lawrence Zitnick. Generative adversarial transformers, 2021. 2, 7
- [27] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks, 2018. 2
- [28] Yifan Jiang, Shiyu Chang, and Zhangyang Wang. Transgan: Two transformers can make one strong gan, 2021. 2, 7
- [29] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs, 2018. 2
- [30] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning, 2016. 7
- [31] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation, 2018. 2, 7
- [32] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data, 2020. 2, 6, 7
- [33] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks, 2019. 2, 8
- [34] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan, 2020. 2, 7

- [35] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning, 2020. [2](#)
- [36] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. [6](#)
- [37] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. [1](#)
- [38] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network, 2017. [2](#)
- [39] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows, 2021. [2](#)
- [40] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. [6](#)
- [41] Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks, 2017. [2](#)
- [42] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks, 2018. [2](#), [7](#)
- [43] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks, 2016. [1](#)
- [44] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans, 2016. [6](#)
- [45] Samarth Sinha, Zhengli Zhao, Anirudh Goyal, Colin Raffel, and Augustus Odena. Top-k training of gans: Improving gan performance by throwing away bad samples, 2020. [7](#)
- [46] Yuxuan Song, Qiwei Ye, Minkai Xu, and Tie-Yan Liu. Discriminator contrastive divergence: Semi-amortized generative modeling by exploring energy of the discriminator, 2020. [7](#)
- [47] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks, 2020. [2](#)
- [48] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision, 2021. [3](#)
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2017. [2](#)
- [50] Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity, 2020. [1](#), [6](#), [7](#)
- [51] Wei Wang, Yuan Sun, and Saman Halgamuge. Improving mmd-gan training with repulsive loss function, 2019. [7](#)
- [52] Jiaheng Wei, Minghao Liu, Jiahao Luo, Qiutong Li, James Davis, and Yang Liu. Peergan: Generative adversarial networks with a competing peer discriminator, 2021. [7](#)
- [53] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers, 2021. [2](#)
- [54] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop, 2016. [6](#)
- [55] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S. Huang. Generative image inpainting with contextual attention, 2018. [1](#), [2](#)
- [56] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks, 2019. [1](#), [7](#)
- [57] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip H. S. Torr, and Li Zhang. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers, 2021. [2](#)
- [58] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks, 2020. [2](#)
- [59] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. Sean: Image synthesis with semantic region-adaptive normalization. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2020. [2](#)
- [60] Zhen Zhu, Zhiliang Xu, Ansheng You, and Xiang Bai. Semantically multi-modal image synthesis, 2020. [2](#)